

Positioning MPLS

This document identifies Multi-Protocol Label Switching (MPLS) technology components, describes their functionality, and illustrates the value they provide in Service Provider environments.

MPLS was initially targeted for Service Provider customers; however, Enterprises have begun to show interest in deploying this technology. This document can apply to large Enterprise customer whose networks resemble Service Provider networks in the following areas:

- Size of the network
- Offer “internal services” to different departments within the Enterprise

MPLS compliments IP technology. It is designed to leverage the intelligence associated with IP Routing, and the Switching paradigm associated with Asynchronous Transfer Mode (ATM). MPLS consists of a Control Plane and a Forwarding Plane. The Control Plane builds what is called a “Forwarding Table,” while the Forwarding Plane forwards packets to the appropriate interface (based on the Forwarding Table).

The efficient design of MPLS uses Labels to encapsulate IP packets. A Forwarding Table lists Label Values, which are each associated with determining the outgoing interface for every network prefix. Cisco IOS Software supports two signaling mechanisms to distribute labels: Label Distribution Protocol (LDP) and Resource Reservation Protocol/Traffic Engineering (RSVP / TE).

MPLS comprises the following major components:

1. *MPLS Virtual Private Networks (VPNs)*—provides MPLS-enabled IP networks for Layer 3 and Layer 2 connectivity. Includes two major components:
 1. Layer 3 VPNs—based on Border Gateway Patrol
 2. Layer 2 VPNs—Any Transport over MPLS (AToM)
2. *MPLS Traffic Engineering (TE)*—provides an increased utilization of network bandwidth inventory and for protection services
3. *MPLS Quality of Service (QoS)*—builds upon existing IP QoS mechanisms, and provides preferential treatment to certain types of traffic, based on a QoS attribute (i.e., MPLS EXP).

MPLS VPNs

Layer 3 VPNs

Layer 3 VPNs or BGP VPNs have been the most widely deployed MPLS technology. They use Virtual Routing instances to create a separate routing table for each subscriber,



and use BGP to establish peering relations and signal the VPN-associated labels with each of the corresponding Provider Edge (PE) routers. This results in a highly scalable implementation, because core (P) routers have no information about the VPNs.

BGP VPNs are useful when subscribers want Layer 3 connectivity, and would prefer to offload their routing overhead to a Service Provider. This ensures that a variety of Layer 2 interfaces can be used on either side of a VPN. For example, Site A can use an Ethernet interface, while Site B uses an ATM interface; however, Sites A and B are part of a single VPN.

It is relatively simple to implement multiple topologies with router filtering, including a Hub & Spoke or Full Mesh:

- *Hub and Spoke*—central site is configured to “learn” all the routes from the remote sites, while the remote sites are restricted to “learn” routes only from the central site.
- *Full Mesh* topologies would result in all the sites having the ability to “learn” or import routes from every other site.

Layer 3 VPNs have been deployed in networks that have as many as—seven hundred PE routers. Service Providers are currently providing up to five hundred VPNs, with each VPN containing as many as one thousand sites. A wide variety of routing protocols are available deploy on the subscriber access link (i.e. CE to PE link). These include Static Routes, BGP, RIP and Open Shortest Path First (OSPF). Most VPNs have been deployed with Static Routes, followed by BGP Routing.

Layer 3 VPNs offer advanced capabilities, including Inter-AS and Carrier Supporting Carrier (CSC). These provide hierarchical VPNs, allowing a Service Provider to provide connectivity across multiple administrative networks. Currently, initial deployments of such functionality are becoming more widespread.

Layer 2 VPNs

Layer 2 VPNs refer to the ability of Service Provider customers to provide Layer 2 circuits over an MPLS-enabled IP backbone. It is important to understand the three major components of Layer 2 VPNs:

1. *Layer 2 Transport over MPLS*—Layer 2 circuit carried transparently over a MPLS enabled IP backbone (also known as AToM)
2. *Virtual Private Wire Services*—the ability to add signaling to AToM, and to features such as auto-discovery of CE devices
3. *Virtual Private LAN Services*—the ability to add a Virtual Switch Instances (VSIs) at the PE routers to provide LAN based services over a MPLS enabled IP backbone

The predominant Layer 2 circuits include Ethernet, ATM, Frame Relay, PPP, and HDLC. AToM and Layer 3 VPNs are based on the same concepts, but AToM uses a directed LDP session to distribute the VC Labels (analogous to BGP VPN label). Consequently, core routers are not required to have the knowledge on a per-subscriber basis, resulting in a very scalable architecture.

Prior to the availability of AToM, Service Providers had to build different networks for providing Layer 2 connectivity. For example, a service provider could be required to build an ATM and a Frame Relay network, resulting in increased operational and capital expenses. Layer 2 VPNs on MPLS now enable service providers to combine these different networks, so they can save significantly in terms of these operational and capital expenses.



Layer 2 VPNs and Layer 3 VPNs can be configured on a single box and can be leveraged for increased revenue streams from subscribers.

Layer 2 and Layer 3 VPNs are complimentary in nature. Over time, demand for Layer 2 VPNs could possibly exceed that of Layer 3 VPNs. However, as enterprise subscribers streamline network requirements, it is foreseeable that this may not be the case.

MPLS Traffic Engineering

MPLS TE was initially envisioned as technology that would enable Service Providers to better utilize the available network bandwidth by using alternate paths (i.e. other than the shortest path). It has evolved to provide multiple benefits, including Connectivity Protection using Fast ReRoute and “Tight QoS”. “Tight QoS” results from using MPLS TE and QoS mechanisms together.

MPLS TE uses IGP, IS-IS and OSPF to flood bandwidth information through a network. It also uses RSVP extensions to distribute labels and constraint-based routing to compute paths in the network. These extensions have been defined in rfc3209.

Service Providers that deploy MPLS TE tend to deploy a full mesh of TE tunnels. This creates a logical mesh, even when a physical topology is not a full mesh. In this environment, Service Providers have noticed additional 40%-50% bandwidth availability from the network. This gain is optimal network usage, which leads to a reduction in capital expenses.

MPLS TE provides Connectivity Protection using Fast ReRoute (FRR). FRR protects primary tunnels by using pre-provisioned backup tunnels. During a failure condition, it takes approximately fifty milliseconds for the primary tunnel to switch over to the backup tunnel. FRR is dependent on Layer 3 protection, unlike SONET or SDH protection that occurs at the interface level. The restoration time is therefore dependent on the number of tunnels and the number of prefixes being switched over. This is a key issue that should be considered while deploying an optimal FRR design.

Internal tests of the Cisco FRR implementation have revealed performance of better than 50 milliseconds; however, the restoration time may be higher, depending upon the configuration. FRR can be used to protect Links, Nodes and the entire LSP Path. Since Path protection implies that the failure notification travels to the headend, restoration times are inherently much slower. Most Service Providers are concerned with local failures, and have found that link failures are more common than node failures.

DiffServ Aware Traffic Engineering is the ability to run TE for different classes of traffic. A Service Provider may decide to operate a set of TE Tunnels that utilize the “sub-pool” for Voice traffic. Further, the Service Provider can ensure that these tunnels use an explicit path, on which the shortest path results in the shortest delay. There may be another set of TE Tunnels that uses the “global pool” for non-voice traffic that is not delay sensitive.

It is important to note that MPLS TE is control plane functionality. When Virtual Leased Line (VLL) solutions are defined, the appropriate QoS mechanisms must be configured (ie: queuing or policing) to meet the bandwidth guarantee. Service Providers are beginning to offer VLL services as voice trunks to connect Central Offices as well as PBXs.

MPLS Quality of Service

MPLS QoS leverages existing IP QoS DiffServ mechanisms by enabling them to work on MPLS paths. Certain extensions, including the ability to set and match on the MPLS EXP bits, have been added; however, the fundamental behavior of the QoS mechanism remains unchanged.

MPLS is fundamentally a tunneling technique, so the QoS mechanism allows for a flexible deployment by “tunneling” the subscriber’s QoS through the QoS policies of the Service Provider.

Suppose that a Service Provider uses EXP value of 6 for voice, and EXP values 4 and 3 for non-voice traffic. It could simultaneously provide transparent services to the following Enterprises subscribers with the following QoS maps:

1. Uses Prec 3 for voice and Prec 2 for non-voice traffic
1. Uses Prec 5 for voice and Prec 4 for non-voice traffic

Offering QoS services within an MPLS VPN has become a very attractive value proposition to many Service Providers, but the extent of the QoS deployment varies between customers. While some have deployed only two classes of service –(voice and non-voice), others deploy as many as five classes:

- Best Effort Data
- Interactive Data (i.e., Telnet)
- Mission Critical Data (ERP applications; i.e., SAP, PeopleSoft)
- Video
- Voice

While the DiffServ architecture defines many as sixty-four traffic classes, it is highly unlikely that such a service offering will be business justified. Imagine the plight of the subscribers when they need to select and offer a QoS for the CEO, or VP. This would present a significant administrative challenge, while diminishing the fundamental value of the service differentiation.

Conclusion

MPLS is emerging as a widely acceptable technology, evidenced by the 100+ customer deployments of Cisco MPLS. It is important to note that MPLS is not a replacement for IP. The IP Control Plane is a fundamental component of MPLS. The ability to add the ATM-like Forwarding Plane makes it extremely attractive to both Service Providers and Enterprises.

Service Providers can reduce their time to profitability by as much as 25% by deploying MPLS VPNs, MPLS QoS and MPLS TE, rather than only providing the vanilla connectivity of VPNs.

To sum it, the fundamental value for Service Providers and Enterprises to deploy an MPLS-enabled IP network is the ability to offer Layer 3 and Layer 2 connectivity and shared services (like DHCP, NAT, etc.) over a single network, with a high degree of optimization and utilization of the available network bandwidth using TE and QoS.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2002, Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0203R) 2102823/ETMG 6/02